

SATURN

(Situational Awareness Tool for Urban Responder Networks)

Heather Zwahlen

Aaron Yahr

Danielle Berven

Michael T. Chan

Maximilian Merfeld

Christine Russ

Jason Thornton

MIT Lincoln Laboratory

Lexington, MA

{heatherz | ayahr | danielle.berven | mchan | max.merfeld
| christine.russ | jason.thornton}@ll.mit.edu

Jalal Mapar

Department of Homeland Security
Science and Technology Directorate,
Washington, DC
Jalal.Mapar@dhs.gov

Abstract— SATURN is a prototype system for the intelligent incorporation of output from surveillance camera networks into an enhanced situational awareness display. It is a web-based, service oriented, open standards platform designed to be accessible to any user with a common browser. SATURN fuses information from an array of sensors including real-time feeds from video cameras. The sensor data is displayed within an intuitive map-based view and is coupled with video analytics algorithms, a chat capability, and collaborative tools for annotation. A principal component of the system is the ability to conduct attribute-based searches for people within live video feeds and for vehicles within archived camera footage. This real-time cueing to events involving people or vehicles of interest provides a potential reduction in manpower and shortened response timeline. SATURN is applicable to a broad set of law enforcement, security, and counterterrorism missions typically addressed by urban responders.

Keywords—video analytics; sensor fusion; video; urban responders

I. INTRODUCTION

Urban authorities have a broad set of missions. Duties vary in both the frequency of occurrence and in the complexity of execution. They include everyday public safety missions such as traffic enforcement as well as special event crowd management. Similarly, they may be relatively straightforward, such as the protection of a critical infrastructure building, or they may require coordination with jurisdictional partners across an extensive physical space, as in the case of disaster management and response.

In order to execute these missions, a number of sensor and information systems are typically utilized. These include databases such as geographic information system (GIS) layers and vehicle registration information, as well as sensor feeds such as asset (person and vehicle) location updates and various communication systems. For example, person location updates may be provided by mobile phones. Also included among

these information systems are video camera systems. Typical urban authorities have hundreds of cameras distributed in both indoor and outdoor environments.

A number of challenges, however, exist in the efficient operation of these existing sensor and information systems. Often times they are procured in isolation from one another. This may be due to limited operation and management budgets, immediate needs, or a lack of impartial technical guidance. The result is stovepipe systems both within a single force and across jurisdictions. Additionally, complex operating environments can result in incomplete or sparse coverage of both RF and visual sensor data. Lastly, the systems tend to generate an unmanageable quantity of data. This is particularly evident for video cameras, where the current forensic technique is often to perform manual review of camera data to gather additional information in the wake of an incident.

The SATURN system arose as a means to address these outstanding challenges. It utilizes a service oriented architecture implemented through an enterprise service bus in order to allow sensor systems to be integrated with ease. To ensure relevance to user concepts of operation (CONOPS), it was developed in conjunction with the Beverly Hills Police Department (BHPD). Through this partnership, SATURN also serves as a testbed for the development and assessment of new technologies such as video analytics [Figure 1].

SATURN was developed through iterative cycles of partner feedback, sensor fusion efforts, and prototype demonstrations. After first understanding the mission needs and operations interplay for sensor and information systems, a sensor fusion effort was undertaken. This effort focused on harnessing existing sensors such as live video and combining them with new technologies within an intelligent web-based display. A prototype demonstration for BHPD was conducted that involved the execution of scripted scenarios during which the developers responded to an incident using SATURN, along with unstructured use of the SATURN system by BHPD

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE JUL 2012		2. REPORT TYPE		3. DATES COVERED 00-00-2012 to 00-00-2012	
4. TITLE AND SUBTITLE SATURN (Situational Awareness Tool for Urban Responder Networks)				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Massachusetts Institute of Technology (MIT), Lincoln Laboratory, 244 Wood Street, Lexington, MA, 02420-9108				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES Presented at the 15th International Conference on Information Fusion held in Singapore on 9-12 July 2012. Sponsored in part by Office of Naval Research and Office of Naval Research Global.					
14. ABSTRACT SATURN is a prototype system for the intelligent incorporation of output from surveillance camera networks into an enhanced situational awareness display. It is a web-based, service oriented, open standards platform designed to be accessible to any user with a common browser. SATURN fuses information from an array of sensors including real-time feeds from video cameras. The sensor data is displayed within an intuitive map-based view and is coupled with video analytics algorithms, a chat capability, and collaborative tools for annotation. A principal component of the system is the ability to conduct attribute-based searches for people within live video feeds and for vehicles within archived camera footage. This real-time cueing to events involving people or vehicles of interest provides a potential reduction in manpower and shortened response timeline. SATURN is applicable to a broad set of law enforcement, security, and counterterrorism missions typically addressed by urban responders.					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 8	19a. NAME OF RESPONSIBLE PERSON
a REPORT unclassified	b ABSTRACT unclassified	c THIS PAGE unclassified			

Sensor and Information Systems

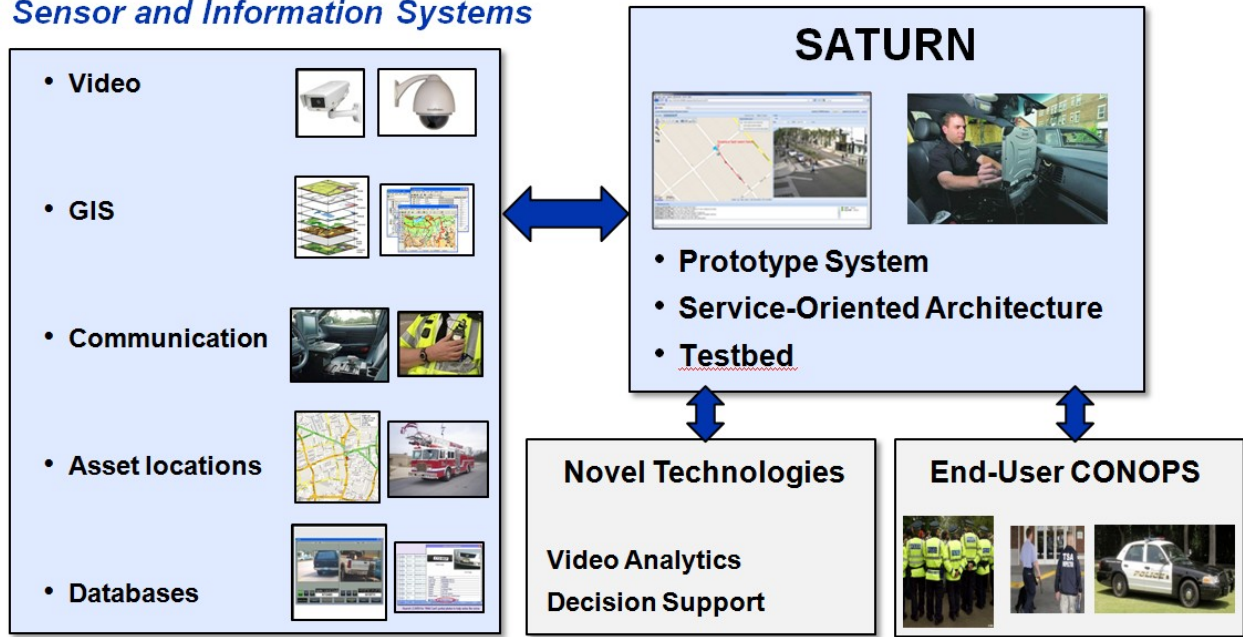


Figure 1. SATURN role

personnel. BPHD users then provided feedback. This in turn informed new sensor integration schemes, system requirements, and CONOPS; and the cycle repeated.

The system fuses information from an array of sensors including real-time feeds from video cameras. The sensor data is displayed in an intuitive map-based view and is coupled with video analytics algorithms, a chat capability, and collaborative tools for annotation.

Various elements of SATURN perform fusion at different levels [1]. This includes person and vehicle detection using multiple cues (Level-1), map-based situation assessment (Level-2), and operator-assisted person and vehicle identification (Level-5), which are all essential to the overall system.

Through these fusion efforts, SATURN serves to:

- Reduce the incident response timeline by changing the search time through surveillance video from hours to seconds when compared to manual review
- Lower the manpower requirement by freeing up human resources for other tasks
- Provide a low cost solution appropriate for the limited operational and management budgets of the typical urban authority

II. KEY SYSTEM ARCHITECTURE COMPONENTS

The SATURN system is based on the service-oriented Next-Generation Incident Command System (NICS) architecture described in [2]. The NICS architecture provides a plug-in architecture and employs a message bus backbone for communication between components. The SATURN project

adds new functionality specifically catered to the urban response environment, including advanced video integration such as intelligent display and analytics capabilities.

A. Service Oriented Architecture (SOA)

In order to address the negative trend of isolated sensor and information systems used by urban authorities, the SATURN backbone adheres to open source principles. An open source methodology eases sensor integration by providing open access to low level software details. The service oriented architecture concept discussed in [2] provides a number of benefits over the current paradigm. By allowing sensors, users, and services to connect in a standardized manner over an enterprise service bus, data fusion and collaboration among users are more easily conducted. In addition, this methodology provides ease in scalability when adding sensors in large quantities, as well as expandability to sensors and services that have not yet been incorporated. Overall, this architecture provides the most efficient use of the sensor and information systems native to local responders.

B. Video Exploitation and Analytics

As mentioned earlier, the ubiquity of surveillance video and its subsequent use by urban authorities cements the need for advanced video integration in order to enable timely use of such video data. Video data has many uses including the relay of situational awareness, recording of incidents in real-time, monitoring of high value assets, and providing cues of suspicious behavior. Although surveillance video tends to capture high-value information content across a range of mission areas, it is often very challenging for operators to extract this information because of the unmanageably large amount of data created by distributed camera systems. Because of limited human resources, manual review of large volumes of surveillance data tends to be a tedious and time-




Figure 2. Attribute-based person search

consuming process. Moreover, it can be difficult for investigators to make sense of the video collected across distributed camera networks because of gaps in the coverage area within complex outdoor and indoor urban environments. This motivates the incorporation of two components that improve situational awareness related to surveillance video: map-based views, and automatic video analysis capabilities.

As a first step in the exploration of an advanced visualization space, SATURN couples local authority camera systems with a map-based view of their locations. This provides an additional means of understanding and visualizing the often disjoint camera networks with which responders must work. Collaborative tools provide the additional means to directly link situational awareness from a specific object or incident of interest to a given camera feed. This allows a more intuitive visualization of the video data and provides new context for the common operating picture.

In addition, video analytics software can help operators make sense of hundreds of cameras feeds spread throughout an urban environment. The purpose of this cutting edge technology is to partially automate the tasks of surveillance video interpretation in order to assist operators with either real-time alerts or forensic analysis. Since responders are often interested in finding specific video content (related to an incident or person or vehicle of interest), the ability to perform automatic content retrieval is a critical need. The following section discusses one novel implementation of content retrieval which has been implemented for SATURN.

III. ATTRIBUTE-BASED SEARCH CAPABILITIES

Video content retrieval can take multiple forms, with queries based on specific actions, motion patterns, or object

types and attributes. The SATURN implementation supports attribute-based queries of persons or vehicles using multiple types of descriptors. First, the software processes all surveillance video, extracts metadata about any observed objects and stores that information to a database. Then, when an operator enters a query, the content retrieval software searches this database for apparent matches and presents them to the user. The major advantage of a system like this is that it enables responders to run fast searches based on vehicle or suspect descriptions and browse the results in order to get to the relevant surveillance data, as opposed to manual scanning very large amounts of video data.

A. Person Search

In order to support attribute-based person search, SATURN implements the novel algorithm described in [3]. This technique is designed to recognize attributes observable at a distance in video, such as clothing type and color, hair color, gender, or the presence of hand-carried objects. It uses a probabilistic model to evaluate the likelihood that each observed pedestrian fits the user-provided description and sorts the results starting with the most likely matches. Figure 2 demonstrates an example of the person-search functionality. In this case, the operator specifies gender and clothing color, and the results (including multiple successful matches) are displayed back to the operator as a set of image chips. This capability is integrated with the other situational awareness components native to the system so that it can be executed on live video for real-time alerts. A similar approach to the problem of searching for vehicles within archived video data was developed in order to expand the usefulness of the search component. This new capability is discussed below.




Figure 3. Attribute-based vehicle search Top left: Result viewing and video playback panel. Top right: Attribute search panel. Bottom right: Detected vehicle and part locations highlighted in red and blue respectively.

B. Vehicle Search

In law enforcement operations, a common example vehicle description received may be “A silver minivan with a white ribbon decal on the back.” In this case, the vehicle attributes are color, type, and a distinguishing feature. Currently the SATURN system supports retrieval of vehicles based on color and type attributes, as described below.

The vehicle search algorithm implemented for SATURN depends on four primary components:

- Vehicle Detector
- Vehicle Type Classifier
- Vehicle Color Classifier
- Background Subtraction

The search algorithm fuses the output of these components to produce ranked search results. The first two components are described in some detail. The last two are based on the same techniques used for the so-called moving person case as in [3].

1) Vehicle Detector

The Vehicle Detector is an algorithm that detects and locates vehicles in images. It is based on the approach in Felzenszwalb et al [4] that uses a part-based mixture model representation, where for example a part can be a rectangular section around a rear tire or the front bumper. The underlying feature is based on Histograms of Oriented Gradients (HoG) [5], but formulated at the part level as well as at the object level. Vehicle models are trained using ground-truthed example

images of vehicles in a variety of viewing conditions and occlusion circumstances. The advantage of using a part-based formulation and a mixture model framework is the ability of the model to gracefully handle partial occlusion and pose variation respectively.

The vehicle models are discriminatively trained using latent Support Vector Machines [6]. The introduction of latent variables to represent qualitative poses and object part locations eliminates the need to label images beyond drawing bounding boxes around the vehicles. The training algorithm is an iterative procedure that alternates between fixing latent values for positive examples and optimizing the latent SVM objective function; details of the overall approach are described in [4].

2) Vehicle Type Classifier

The Vehicle Type Classifier uses the same underlying model representation as the Vehicle Detector. However, the classification models are trained differently. Given N vehicle types of interest, each of the corresponding N classifiers are trained using image chips of the target vehicle type as positive training samples, and image chips of the other $N-1$ types as negative training samples. The bounding box location of the negative samples (as well as positive ones) is also treated as a latent variable, which is a distinct difference in how the Vehicle Detector is trained. The training process starts by substituting the negative samples with random non-vehicle images to improve convergence.

3) Training Data and Model Learning

There are three sources of image data for the vehicle models. The first data source was the VOC2010 database that is publicly available from the PASCAL community [7]. The second was archival data from an existing law enforcement

video surveillance system. The third was online data from autotrader.com.

Vehicles are divided into eight common types: Sedan, Coupe, Convertible, Hatchback, Station Wagon, Van, SUV, and Pickup Truck. The choice to use these eight common types was largely based on classification schemes used by many online classifieds web sites (such as autotrader.com), which reflects how people commonly describe vehicles. Vehicles are also characterized by 15 perceptual colors commonly used for describing cars in classified ads.

A set of MATLAB-based tools was developed to establish the ground truth for a subset of the data. The tools allow one to mark bounding boxes around vehicles in video (and static images), specify their types and primary colors, and select representative points on vehicles with the specified color.

The type models and color models were developed separately as they are largely independent. Multivariate Gaussian distribution in HSV space was used for the color models, and latent-SVM for the type classification models described in the previous section.

4) Vehicle Retrieval in Video

The vehicle retrieval algorithm for moving vehicles, given a specified color and type, follows these steps. First, the Vehicle Detector is applied to detect vehicles in video at a rate of once every second. An associated foreground motion score is also calculated in each of the corresponding detected regions based on a statistical background subtraction process [8]. Only those candidates with detection scores and motion scores above certain thresholds are kept. The eight Vehicle Type Classifiers are then applied to the detections, producing eight type scores for each of them, and the 15 color likelihood scores are calculated based on the trained color models. The detection bounding boxes, times of detection, and the associated scores are then stored in a database.

When performing a search for a vehicle of a particular type and color (usually within a time window), a fused score defined as a weighted sum of the color likelihood score and a z-normalized SVM classification score is calculated for each potential detection. The detections are then rank ordered by their fused scores. In practice, to reduce multiple instances of any particular vehicle being returned, non-maximal suppression was performed by reducing the scores of the detections within a small time window (Gaussian shape of about 10 seconds wide).

This is done first for the top-ranked detection, and then recursively for the successively lower ranked candidates.

The top matched candidates (typically 30) are presented to the operator through the SATURN system in reverse rank order. The operator then tries to visually identify the vehicle of interest from this short list based on other attributes that human can distinguish more readily (e.g. the presence of a decorative decal). Figure 3 illustrates the vehicle retrieval capability.

The accuracy of the algorithm has not yet been established through precision/recall performance analysis. However, observations indicate some vehicle types are more discriminative than others. For example, the Vehicle Type Classifiers have equal error rates ranging from about 20% (for pickup trucks) to 35% (for station wagon) based on a validation data subset that was not used in training; the actual performance is likely better as the ultimate models used in our system were trained with all of the training and validation data.

There are some remaining technical challenges for the algorithm. At the training phase, one obstacle included obtaining a balanced data set. Examples of vehicles such as sedans and SUVs tended to be more common in the available video than, for example, station wagons. At the classification phase, additional work is needed to determine how to best leverage the confusion matrix between the eight vehicle types. One possible solution is a cascading approach to classification. Finally, the scoring metric weighting color and type matches may be further optimized. Prior knowledge of the vehicle color distribution (e.g. rarity of lime green) suggests using unique weights for each color type. User defined input based on confidence of witness description is another option. Future effort will further investigate these areas.

IV. SATURN SYSTEM

The current SATURN system is the result of iterative prototype demonstrations for the Beverly Hills Police Department. Building on the NICS architecture as described in [2], the system further integrates live feeds from city video cameras, geo-location feedback from both responder vehicles and individuals to the map display, and attribute-based person and vehicle searches based on camera streams. The user interface and overall sensor integration effort are detailed below.

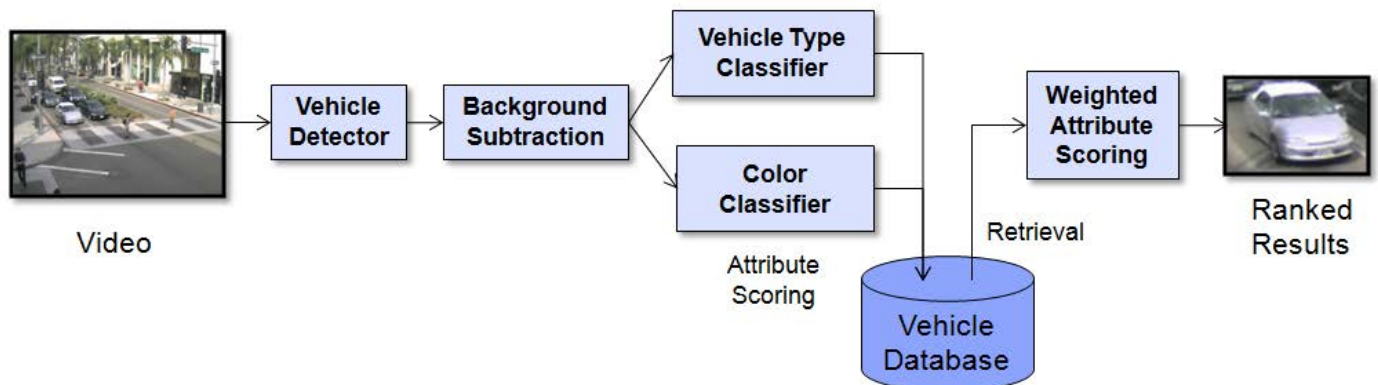


Figure 4. Attribute-based vehicle search

A. Graphical User Interface (GUI)

The SATURN graphical user interface (GUI) features the mapping, white board, and chat functionality described in [2]. An “incident” may be created at the start of an investigation which allows collaboration such as white board markings and text chat for a specific user set. Multiple, simultaneous incidents are supported, and all incident data is archived for future viewing. SATURN incident types have been customized for law enforcement and include categories such as robbery or assault.

Depending on the camera chosen, both live and archived video streams may be viewed. Additionally, attribute-based searches may be performed on the video data. Vehicle searches are currently implemented only on archival camera data for demonstration purposes; however searches on live video could be incorporated into a future operational system. Person searches are available for both live video and archival video. Top results are displayed along with a likelihood score. Searches conducted in the live mode continually update the list of top matches every ten seconds with the newest high-ranking detections. In this manner, the user is able to receive real-time alerts to a person of interest fitting a prescribed description.

B. Sensor Integration and Fusion

SATURN leverages the map-based view and collaborative tools to provide further advanced video integration for an urban authority’s camera system. Camera locations are denoted on the map by clickable icons. Selecting an icon opens a tab adjacent to the map view to provide the user with immediate viewing capability.

Upon selection of a specific camera, the user may quickly transition to the input tab for attribute-based searches. Once a search has been conducted, results appear on the map view and are denoted by either car or person icons. If conducted with an incident, search results may be further shared among collaborators for quick dissemination of relevant information. Selecting a given result automatically cues the video to the

point of detection and highlights the object of interest within the scene.

Finally, the position of mobile assets is available in real-time. The updates of both urban responder personnel and their vehicle locations are displayed on the map view to indicate their proximity to an incident or object of interest.

Figure 5 illustrates relevant features of the SATURN web GUI. In this example, an attribute-based person search has been performed on video data. Results of interest are denoted on the map view by the blue person icons, and the original search inputs are shown on the right. In this manner, the SATURN system is able to provide advanced video integration in the form of both the intelligent geo-referenced display of video metadata and content retrieval capabilities to reduce the resources required to forensically parse camera scenes. Overall, these improvements allow real-time alerts to objects of interest as part of a more complete situational awareness picture.

In the case of the Beverly Hills Police Department, the system may aid in both everyday police work such as locating stolen cars, or in less frequent duties such as crowd management at special events. The alerts provided through the video analytics capability help push the use of surveillance video in incident response from a forensic, after-event use to a real time aid.

V. TECHNICAL DETAILS

A. Video Viewing and Processing

Figure 6 provides an overview of the implementation details. Archived video is accessed using HTML5 on the client side. The HTML5 standard is the first version of HTML to incorporate timed multimedia playback and streaming, thus allowing browsers to support audio and video playback absent a third party plug-in. The camera stream is also transcoded into a flash stream accessible through an HTTP video server

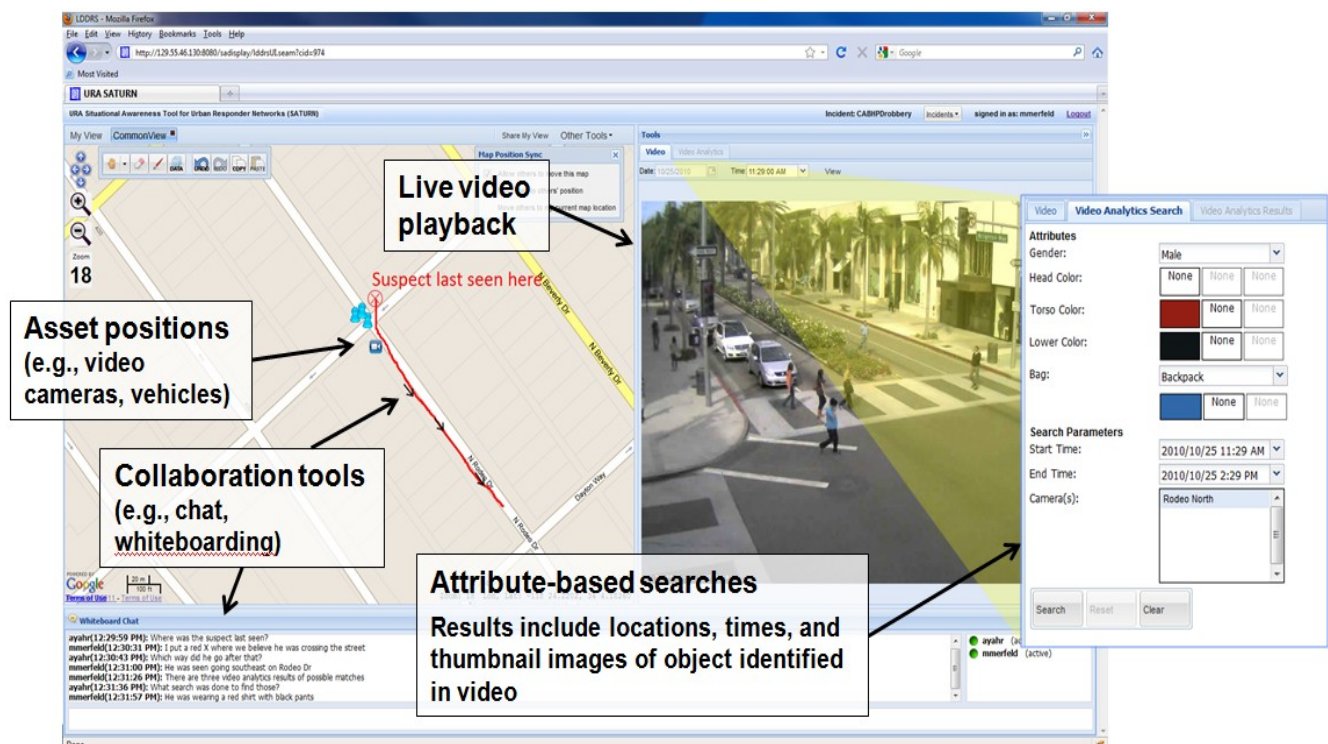


Figure 5. SATURN GUI

deployed on the SATURN server to allow for viewing of the live video. The client-side browser requires a standard Adobe Flash plug-in to view the live video.

The IP cameras leveraged in the SATURN system are part of Beverly Hills City's private network and are not directly accessible to the outside world. The SATURN server requires a VPN connection to the Beverly Hills network to create a HTTP connection to the IP-based cameras. The SATURN server runs an Apache [9] HTTP server that forwards all camera stream requests through the VPN connection to the correct IP address.

The VPN Listener component is responsible for processing the video stream and sending it to the VA processor as well as storing video clips. After an HTTP connection to an IP camera of interest is established, the incoming video is transcoded to correct the bit rate and format so that it is accessible to the remote client and stored on a network file system (NFS) on the SATURN server. One frame per second of raw imagery is sent to the VA server through a UDP connection. The video is stored in the video clip data store as a series of clips of predefined length and the relative metadata is stored in a database, so that the correct video clip can be found easily. The video is stored at 600 kbps, a rate chosen to reflect the resolution required for mission needs. Illustrative Verizon 4G network measurements were taken throughout the city of Beverly Hills in early 2011. Approximately 90% of the time, ample bandwidth of around 2 Mbps was available to meet streaming video requirements.

Real-time video analytics searches are accomplished on live video using a cascaded processing approach shown in Figure 5. The VA Listener receives one frame per second from the VPN Listener. The image frame is placed in a shared memory buffer so that it can be quickly accessed by the Histogram of

Gradients Classification [5] process and then the Results Filtering and Database creation process. In order to achieve less than one second of latency per frame, the Person Detection and Classification process is deployed on an NVIDIA Tesla video processor card using Dalley's GPU implementation of parallelized HOG feature computation [10]. The results filtering and database creation process filters results based on ground-plane and foreground motion information as described in [3] and stores the results.

The GUI supports time-windowed video searches for people and vehicles based on their attributes. Person searches are requested by the web clients based on attributes such as clothing and bag color. Vehicle searches are based on vehicle type and color. For both types of searches, the Query Adapter retrieves results from both archived video analysis as well as live video streams. The results of the query are sent via the message bus to the Image Chip Generator which creates image chips highlighting the search results. The web client receives the results message and retrieves the image chips for display.

B. Hardware Overview

The hardware implementation of the SATURN server architecture consists of two physical machines that reside on the MIT Lincoln Laboratory network. The first, the SATURN server, is dedicated to obtaining and storing the various data feeds as well as running the application server that hosts the SATURN web application. The second, the video analytics (VA) server, runs all of the live VA processing and contains the database of person and vehicle detections. The two machines are connected to the same subnet on the Lincoln Laboratory network and communicate via UDP transmissions as well as across the enterprise message bus. This is implemented using

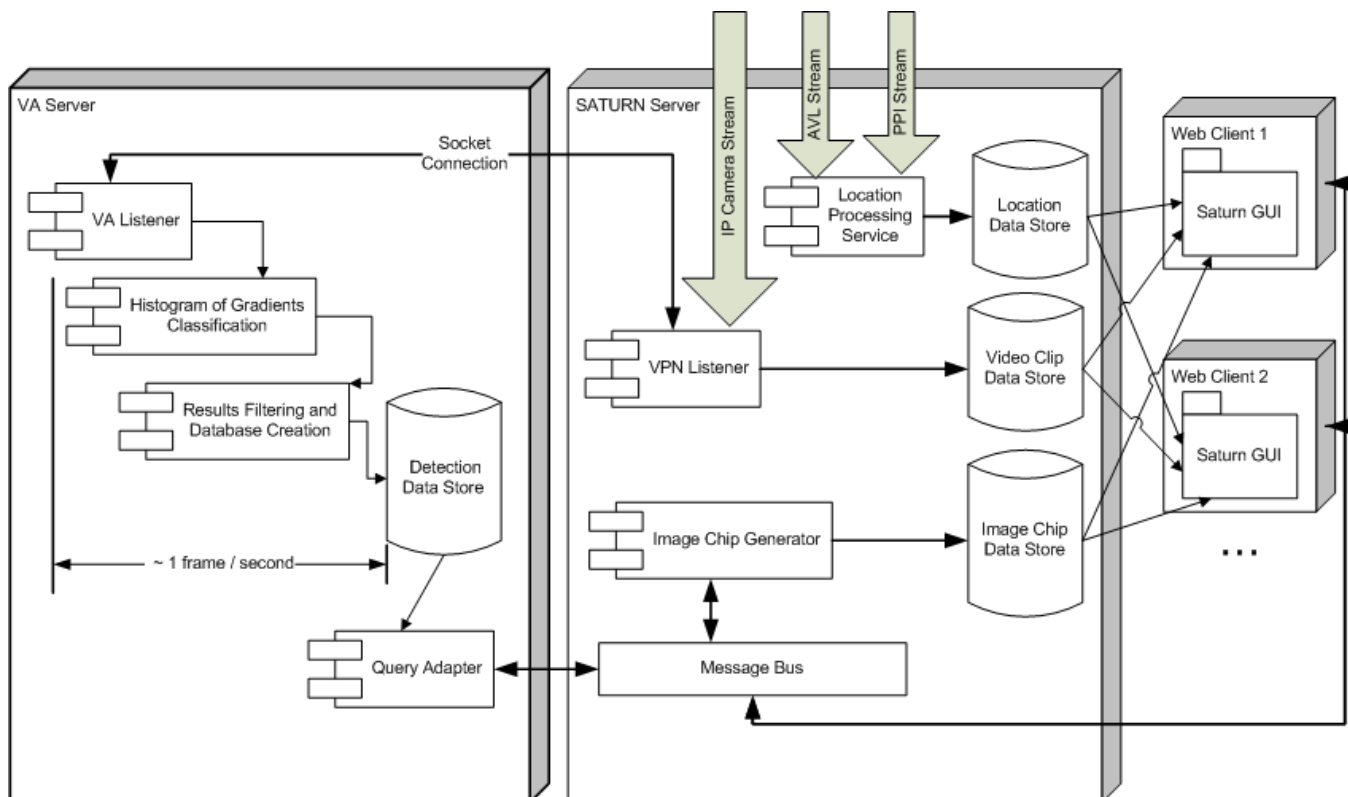


Figure 6. Technical Details

RabbitMQ messaging based on the Advanced Message Queuing Protocol (AMQP) [11].

The SATURN server is a single physical machine that hosts a collection of virtual machines (VMs). The use of VMs allows the use multiple operating systems and makes hardware resource allocation to specific functions easier. The VMs do not share physical memory and therefore must communicate across the network. Currently, there are three VMs on the web server: the deployment VM, the VPN (Virtual Private Network) VM, and the video processing VM. The deployment VM hosts the SATURN application server and stores all available content. The VPN VM connects to the Beverly Hills network through a VPN tunnel and forwards the IP camera feeds to the video processing VM. The video processing VM does all of the video processing, including transcoding and storing the video for client viewing, as well as sending static images through UDP messages to the video analytics machine.

C. Location Processing

In addition to video processing, the SATURN system supports location data streams relevant to law enforcement applications. Two different types of location streams are currently obtained: automatic vehicle location (AVL) streams and mobile personal position indicator (PPI) streams. The AVL streams are sent from AVL units located on certain Beverly Hills fire trucks and police vehicles. The city of Beverly Hills has a large number of first responder vehicles outfitted with AVL units and has programmed a subset of them to forward data to the URA servers. The mobile PPI streams come from a set of smart phones that were programmed to send their location information at periodic intervals. The location processing component parses the data into a set of key markup language (KML) files which are stored in the location data store.

VI. CONCLUSIONS

The SATURN system fuses existing sensor systems of local responders to provide a situational awareness platform coupled with advanced video integration. Through iterative prototype demonstrations with the Beverly Hills Police Department, SATURN served as a testbed for the development of a novel video analytics attribute-based vehicle search capability by incorporating realistic mission needs and CONOPS. Similarly, an existing attribute-based person search algorithm was implemented for the first time on live video to provide real-time alerts.

Future work will continue to develop system enhancements with the objective of producing actionable situational awareness derived from the unification of urban authorities' sensor and information systems. This will include development of a prototype scaled to take in larger numbers of sensors, for example hundreds of video cameras. Through the development cycle, novel technologies such as new analytics techniques and ad-hoc networking in disadvantaged communications environments will continue to be explored.

Lastly, one new effort is expected to focus on advanced visualization that will include improved handling of large video systems, mobile displays, and a virtual command center.

These pieces will help build the foundation required to tackle the needs of higher complexity missions such as disaster response management. In scenarios such as this, many sensor systems and data feeds will need to be interoperable by users across multiple federal and local jurisdictions. In this manner, SATURN will continue to make significant, holistic contributions to information fusion encompassing a wide range of sensor and information systems such as communications, video, GIS layers, maps, and databases. These innovations promise to be the foundation for developing new information sharing technologies for the homeland security user community.

ACKNOWLEDGMENTS

The authors would like to thank the members of the Beverly Hills Police Department for their input and support. This work is sponsored by the Department of Homeland Security Science and Technology Directorate under Air Force Contract FA8721-05-C-0002. Opinions, interpretations, conclusions, and recommendations are those of the authors and are not necessarily endorsed by the United States Government.

REFERENCES

- [1] F. White, Jr., Data Fusion Lexicon, Joint Directors of Laboratories, Technical Panel for C3, Data Fusion Sub-Panel, Naval Ocean Systems Center, San Diego, 1987.
- [2] R. Di Ciaccio, J. Pullen, P. Breimyer, "Enabling distributed command and control with standards-based geospatial collaboration," *Technologies for Homeland Security (HST), 2011 IEEE International Conference on*, vol., no., pp.512-517, 15-17 Nov. 2011.
- [3] J. Thornton, J. Baran-Gale, D. Butler, M. Chan, H. Zwahlen, "Person attribute search for large-area video surveillance," *Technologies for Homeland Security (HST), 2011 IEEE International Conference on*, vol., no., pp.55-61, 15-17 Nov. 2011.
- [4] P. Felzenszwalb, R. Girshick, D. McAllester, D. Ramanan, "Object Detection with Discriminatively Trained Part Based Models," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 32, No. 9, September 2010.
- [5] N. Dalal and B. Triggs, "Histograms of oriented gradients for human detection," in *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, 2005.
- [6] S. Andrews, I. Tschantzaris, and T. Hofmann, "Support vector machines for multiple-instance learning," in *Advances in Neural Information Processing Systems*, 2003.
- [7] M. Everingham, L. Van Gool, C. Williams, J. Winn, A. Zisserman, "The PASCAL Visual Object Classes Challenge 2010 (VOC2010) Results," <http://www.pascal-network.org/challenges/VOC/voc2010/workshop/index.html>.
- [8] C. Stauffer and W. Grimson, "Adaptive background mixture models for real-time tracking," *CVPR*, pp. 246-52, 1999.
- [9] The Apache software foundation, [Online], Available: <http://apache.org>.
- [10] G. Dalley, "Improved robustness and efficiency for automatic visual site monitoring," Ph.D. dissertation, Massachusetts Institute of Technology, Dept. of Electrical Engineering and Computer Science, 2009.
- [11] RabbitMQ, [Online], Available: <http://www.rabbitmq.com>.